Fast Localization of Facial Landmark Points

Nenad Markuš*, Miroslav Frljak*, Igor S. Pandžić*, Jörgen Ahlberg†, and Robert Forchheimer†

* University of Zagreb, Faculty of Electrical Engineering and Computing, Unska 3, 10000 Zagreb, Croatia
† Linköping University, Department of Electrical Engineering, SE-581 83 Linköping, Sweden

March 28, 2014

arXiv:1403.6888v1 [cs.CV] 26 Mar 2014

Abstract

The main objective of this work is to develop a method that can accurately estimate the positions of relevant facial landmarks in real-time on hardware with limited processing power, such as modern mobile devices. This is achieved with a sequence of estimators based on ensembles of regression trees. The trees use simple pixel intensity comparisons in their internal nodes and this makes them able to process image regions very fast. We test the developed system on several publicly available datasets and analyse its processing speed on various devices.

1 Introduction

The human face is central to our identity. It plays an essential role in everyday interaction, communication and other routine activities. Thus, automatic face analysis systems potentially open a wide range of applications. Using the face as a means of human-computer interaction is helping disabled people improve their daily lives, and may become a hands-free alternative in other applications or an entertaining element in innovative games. Similarly, tracking facial actions is a basis for driving computer animated faces in games and entertainment applications. In display technology, view-dependent rendering relies on head tracking to dynamically generate the correct perspective on a 3D scene depending on the location of the user, creating a sensation of 3D on regular 2D screens. In security and authentication applications, face analysis is a front-end for facial recognition and may also be used for liveness detection to avoid fraud by substituting the live face with an image. In augmented reality applications, face tracking is a basis for augmenting the face with additional graphics, enabling commercial applications such as virtual try-on of eyewear, hairstyles or makeup, as well as games and artistic creations. Facial tracking is increasingly finding use in safety applications to detect situations such as sleepiness or lack of attention while driving or using hazardous machinery. There is also a growing interest in automatic processing of social signals [20], which often involves face analysis. Model-based coding of facial video relies on facial landmark point tracking to enable very low bitrate video communication. Last but not least, facial tracking is a tool for analysing facial motion in various research fields, such as chewing analysis within food-related research. All these applications create a strong motivation for the development of fast and accurate automatic face analysis systems.

In this paper we concentrate on the problem of automatic estimation of salient facial landmark positions, such as eye corners or the tip of the nose, from an image region containing a face. This is known in the literature as *face alignment*, *facial landmark point localization*, or *face parts localization*.

1.1 Relevant recent work

Significant progress has been achieved recently in the area of facial landmark localization. The methods considered to be state-of-the-art are described in [4, 5, 18]. The approach described by Belhumeur et al. [4] outperformed previously reported work by a large margin. It combines the output of SIFT-based face part detectors with a non-parametric global shape model for the part locations. The main drawback with this approach is its low processing speed. Cao et al. [5] described a regression-based method for face alignment. Their idea is to learn a function that directly maps the whole facial shape from the image as a single high-dimensional vector. The inherent shape constraint is naturally encoded in the output. This makes it possible to avoid parametric shape models commonly used by previous methods. As this is a tough regression problem, they fit the shape in a coarse-to-fine manner using a sequence of fern[1] ensembles with shape-indexed pixel intensity comparison features. The developed system is both fast and accurate. The system developed by Sun et al. [18] is based on a deep convolutional neural network trained to estimate the positions of five facial landmarks. Additionally, simpler networks are used to further refine the results. The authors report state-of-the-art accuracy results. Recently, deep neural networks started to outperform other methods on many machine learning benchmarks (for example, see [10]). Thus,

[1] A simplified decision tree, see [16].

this is not at all surprising. However, neural networks are usually slow at runtime as they require a lot of floating point computations to produce their output, which is particularly problematic on mobile devices. Chevallier et al. [6] described a method similar to the one we present in this paper. We address this later in the text.

1.2 About our work

In this paper we present results obtained by extending our previous work in eye center localization [14] to the problem of face alignment. The developed prototype achieves encouraging accuracy results and runs in real-time on hardware with limited processing power, such as mobile devices. This indicates that the method can be used to perform fast facial feature tracking on these devices.

2 Method

The basic idea is to use a multi-scale sequence of regression tree-based estimators to infer the position of each facial landmark point within a given face region. We assume that this region is known in advance. This does not pose a problem since very efficient and accurate face detectors exist (including our own work, currently submitted for review [13]). In contrast to most prior work, we treat each landmark point separately, disregarding the correlation between their positions. Of course, a shape constraint could be enforced in the post-processing step and there are many methods to achieve this. We have decided to exclude this step in order to focus on landmark localization itself. We explain the details of the method in the rest of this section and compare it with previous approaches.

2.1 Regression trees based on pixel intensity comparisons

To address the problem of image based regression, we use an optimized binary decision tree with pixel intensity comparisons as binary tests in its internal nodes. Variations of this approach have already been used by other researchers to solve certain computer vision problems (for example, see [17, 12, 16]). We define a pixel intensity comparison binary test on image I as

$$\text{bintest}(I; \mathbf{l}_1, \mathbf{l}_2) = \begin{cases} 0, & I(\mathbf{l}_1) \leq I(\mathbf{l}_2) \\ 1, & \text{otherwise,} \end{cases}$$

where $I(\mathbf{l}_i)$ is the pixel intensity at location $\mathbf{l}_i$. Locations $\mathbf{l}_1$ and $\mathbf{l}_2$ are in normalized coordinates, i.e., both are from the set $[-1, +1] \times [-1, +1]$. This means that the binary tests can easily be resized if needed. Each terminal node of the tree contains a vector that models the output. In our case, this vector is two-dimensional since we are interested in estimating the landmark position within a given image region.

The construction of the tree is supervised. The training set consists of images annotated with values in $\mathbb{R}^2$. In our case, these values represent the location of the landmark point in normalized coordinates. The parameters of each binary test in internal nodes of the tree are optimized in a way to maximize clustering quality obtained when the incoming training data is split by the test. This is performed by minimizing

$$Q = \sum_{\mathbf{x} \in C_0} \|\mathbf{x} - \bar{\mathbf{x}}_0\|_2^2 + \sum_{\mathbf{x} \in C_1} \|\mathbf{x} - \bar{\mathbf{x}}_1\|_2^2, \tag{1}$$

where C_0 and C_1 are clusters that contain landmark point coordinates $\mathbf{x} \in \mathbb{R}^2$ of all face regions for which the outputs of binary test were 0 and 1, respectively. The vector $\bar{\mathbf{x}}_0$ is the mean of C_0 and $\bar{\mathbf{x}}_1$ is the mean of C_1. As the set of all pixel intensity comparisons is prohibitively large, we generate only a small subset[2] during optimization of each internal node by repeated sampling of two locations from a uniform distribution on a square $[-1, +1] \times [-1, +1]$. The test that achieves the smallest error according to equation 1 is selected. The training data is recursively clustered in this fashion until some termination condition is met. In our setup, we limit the depth of our trees to reduce training time, runtime processing speed and memory requirements. The output value associated with each terminal node is obtained as the weighted average of ground truths that arrived there during training.

It is well known that a single tree will most likely overfit the training data. On the other hand, an ensemble of trees can achieve impressive results. A popular way of combining multiple trees is the gradient boosting procedure [9]. The basic idea is to grow trees sequentially. Each new one added to the ensemble is learned to reduce the remaining training error further. Its output is shrunk by a scalar factor called shrinkage, a real number in the set $(0, 1]$, that plays a role similar to the learning rate in neural networks training. We set this value to 0.5 in our experiments.

2.2 Estimating the position of a landmark point

We have observed that accuracy and robustness of the method critically depend on the scale of the rectangle within which we perform the estimation. If the rectangle is too small, we risk that it will not contain the facial landmark at all due to the uncertainty introduced by face tracker/detector used to localize the rectangle. If the rectangle is big, the detection is more robust but accuracy suffers. To minimize these effects, we learn multiple tree ensembles, each for estimation at different scale. The method proceeds in a recursive manner, starting with an ensemble learned for largest scale. The obtained intermediate result is used to position the rectangle for

[2] 128 in our implementation.

the next ensemble in the chain. The process continues until the last one is reached. Its output is accepted as the final result. This was inspired by the work done by Ong et al. [15].

The output of regression trees is noisy and can be unreliable in some frames, especially if the video stream is supplied from a low quality camera. This can be attributed to variance of the regressor as well as to the simplicity of binary test at internal nodes of the trees: Pixel footprint size changes significantly with variations in scale of the eyes and we can expect problems with aliasing and random noise. These problems can be reduced during runtime with random perturbations [8]. The idea is to sample multiple rectangular regions at different positions and scales around the face and estimate the landmark point position in each of them. We obtain the result as the median over the estimations for each spatial dimension.

We would like to note that Chevallier et al. [6] described a similar method for face alignment. The main difference is that they use Haar-like features instead of pixel intensity comparisons to form binary tests in internal nodes of the trees. Also, they do not perform random perturbations in runtime. This is presumably not needed with Haar-like features as they are based on region averaging which is equivalent to low pass filtering and this makes them more robust to aliasing and noise.

3 Experimental analysis

We are interested in evaluating the usefulness of the method in relevant applications. Thus, we provide experimental analysis of its implementation in the C programming language. We compare its accuracy with the reported state-of-the-art and modern commercial software. Also, we analyse its processing speed and memory requirements.

3.1 Learning the estimation structures

We use the LFW dataset [7] and the one provided by Visage Technologies (http://www.visagetechnologies.com/). Both consist of face images with annotated coordinates of facial landmarks. These include the locations of eyebrows, nose, upper and lower lip, mouth and eye corners. Overall, the total number of annotated faces is around 15 000. We intentionally introduce position and scale perturbations in the training data in order to make our system more robust. We extract a number of samples from each image by randomly perturbing the bounding box of the face. Furthermore, as faces are symmetric, we double the size of the training data by mirroring the images and modifying the landmark point coordinates in an appropriate manner. This process results in a training set that consists of approximately 10 000 000 samples.

Each landmark point position estimation structure is learned independently in our framework. We have empirically found that 6 stages with 20 trees of depth equal to 9 give good results in practice. The ensemble of the first stage is learned to estimate the position of a particular landmark point from the bounding box of the face. Each next stage is learned on a training set generated by shrinking the bounding box by 0.7 and repositioning its center at the position output by the ensemble of the previous stage. This process proceeds until the last stage is reached. The learning of the whole estimation structure for a single landmark point takes around one day on a modern desktop computer with 4 cores and 16 GB of memory.

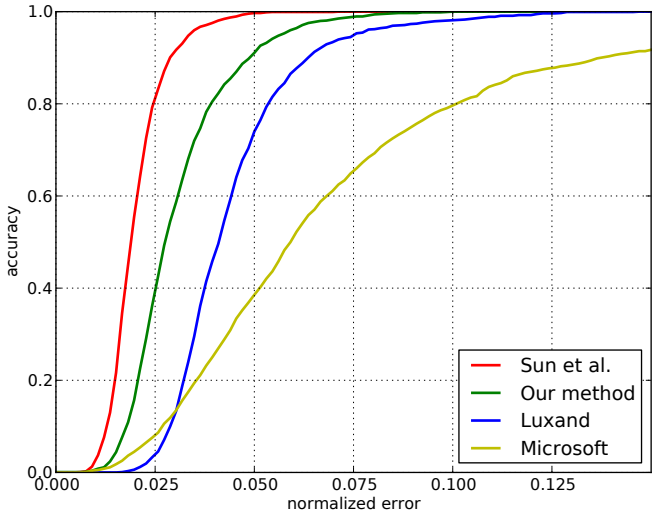

Figure 1: Accuracy curves on the BioID dataset.

3.2 Accuracy analysis on still images

We use the BioID [11] and LFPW [4] face datasets to evaluate the accuracy in still images. The normalized error [11] is adopted as the accuracy measure for the estimated landmark point locations:

$$e = \frac{1}{N} \sum_{n=1}^{N} \frac{D_n}{D}, \tag{2}$$

where N is the number of facial landmarks, D is the distance between the eyes and D_n is the Euclidean distance between the estimated landmark position and the ground truth. The accuracy is defined as the fraction of the estimates having an error smaller than a given number. Roughly, an error of 0.25 corresponds to the distance between the eye center and the eye corners, 0.1 corresponds to the diameter of the iris, and 0.05 corresponds to the diameter of the pupil.

First, we use the BioID and LFPW datasets to compare our system to the state-of-the-art method based on convolutional neural networks [18] and two modern commercial systems, one provided by Microsoft [2] and the other by Luxand [1]. We follow the protocol from [18] to obtain normalized errors for five facial landmarks (eye centers, mouth corners and tip of the nose). The results are reported in figures 1 and 2. We

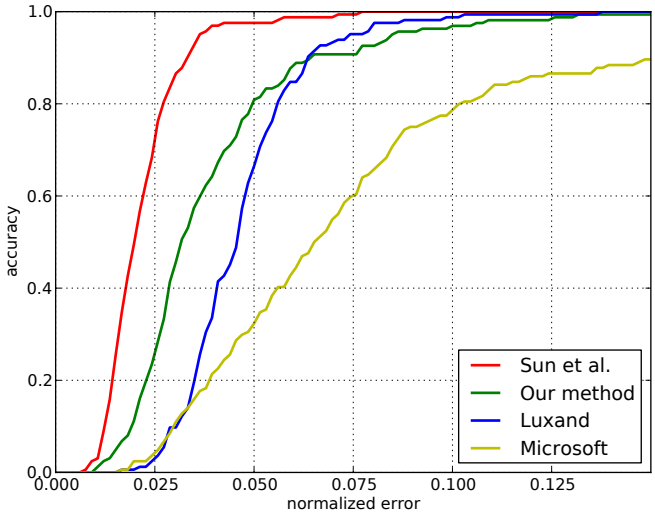

Figure 2: Accuracy curves on the LFPW dataset.

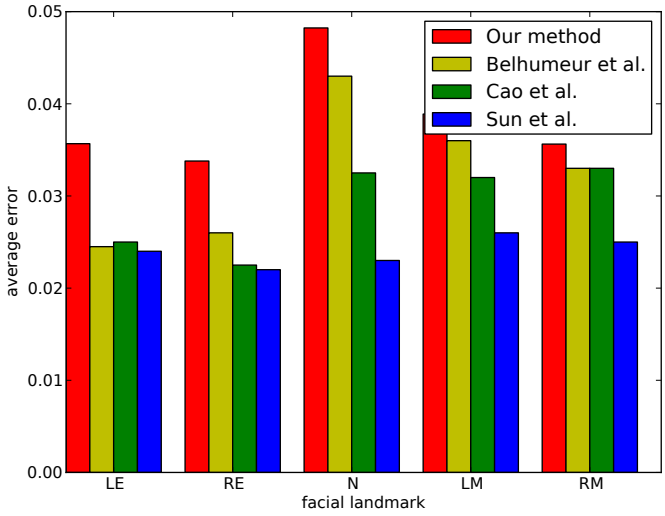

Figure 3: Average accuracy comparison on the LFPW test set (249 images).

can see that our method outperforms both commercial systems in accurate landmark point localization ($e \approx 0.05$) but the neural network-based system is clearly the best. We could not include the method by Cao et al. [5] in this comparison since no code/binaries are available to reproduce the results. We also excluded the results published by Chevallier et al. [6] since they used the evaluation methodology where they partitioned BioID in two parts: one for training and the other for accuracy analysis. It has been argued that this evaluation methodology is flawed since the learning procedure overfits some particular features present only in the used dataset and thus yields performance that is not representative in the general case [19].

In order to compare our method also with the methods excluded from the first experiment, we performed a second comparison. This one is based on average errors reported on the LFPW dataset (accuracy curves could not be obtained due to the lack of data in [4] and [5]). The average error for 5 facial landmarks can be seen in Figure 3. As the average error is sensitive to outliers and LFPW faces vary greatly in pose and occlusions, our method performs worse than other approaches that use some form of shape constraint. Nevertheless, we can see that, on average, the landmark positions estimated by our system are within the pupil diameter from the ground truth.

3.3 Tracking facial features

We use the Talking Face Video [3] to evaluate our system quantitatively in real-time applications. The video contains 5000 frames taken from a video of a person engaged in conversation. A number of facial landmarks were annotated semi-automatically for each frame with an active appearance model trained specifically for the person in the video. These annotations include the locations of eye centers, mouth corners and the tip of the nose. The normalized error averaged over the video sequence obtained by our system was equal to 0.028. Accuracy curve can be seen in Figure 4. These results show that most of the time our system estimated the positions of facial landmarks with high accuracy.

Device	CPU	Time [ms]
PC1	3.4GHz Core i7-2600	1.3
PC2	2.53GHz Core 2 Duo P8700	2.1
iPhone 5	1.3GHz Apple A6	5.0
iPad 2	1GHz ARM Cortex-A9	8.8
iPhone 4S	800MHz ARM Cortex-A9	10.9

Table 1: Average times required to align 5 facial landmarks.

3.4 Processing speed and memory requirements

Processing speeds obtained by our system on various devices can be seen in Table 1. Our system uses a single CPU core although the computations can easily be parallelized. Both Cao et al. [5] and Sun et al. [18] vaguely[3] report processing speeds on modern CPUs: their systems localize 29 and 5 facial landmarks, respectively, in 5 and 120 [ms], respectively.

Each landmark position estimator consisting of 120 trees, each of depth 9, requires around 700 kB of memory. In our opinion, these relatively large memory requirements are one of the drawbacks with our approach as they are inconvenient for some applications, such as face tracking in web browsers or on mobile devices. The problem can be addressed by quantizing

[3]For example, we are not sure if they used multi-core processing in runtime (both papers mention it at some point).

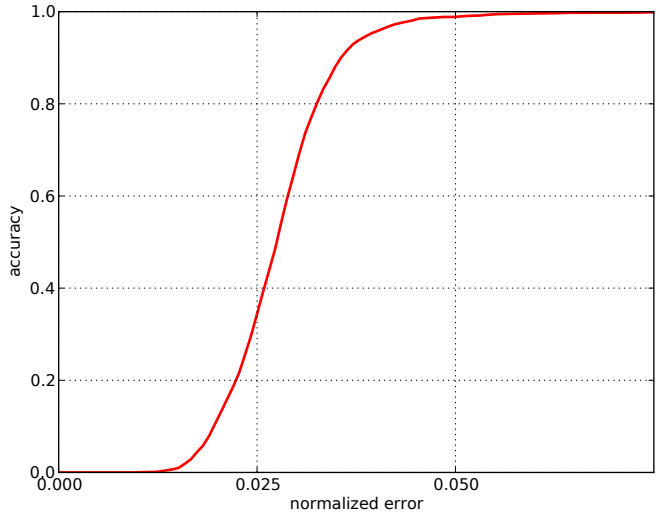

Figure 4: Accuracy curve for The Talking Face Video.

the outputs in the leafs of each tree. In the current implementation, we represent each output with two 32-bit floating point values.

3.5 Qualitative results

Some qualitative results obtained by our system can be seen in Figure 5 and in the video available at http://youtu.be/xpBXpI39s9c. Furthermore, we prepared a demo application for readers who wish to test the method themselves. It is available at http://public.tel.fer.hr/lploc/.

4 Conclusion

Numerical results show that our system is less accurate than the reported state-of-the-art but more accurate than two modern commercial products while being considerably faster, in some cases even by a factor of 50. Its landmark point position estimations are, on average, in the pupil diameter ($e \approx 0.05$) from human-annotated ground truth values. Processing speed analysis shows that the system can run in real-time on hardware with limited processing power, such as modern mobile devices. This enables fast and reasonably accurate facial feature tracking on these devices. We believe that the method described in this paper achieves acceptable accuracy and processing speed for a lot of practical applications.

Acknowledgements

This research is partially supported by Visage Technologies AB, Linköping, Sweden, by the Ministry of Science, Education and Sports of the Republic of Croatia, grant number 036-0362027-2028 "Embodied Conversational Agents for Services in Networked and Mobile Environments" and by the European Union through ACROSS project, 285939 FP7-REGPOT-2011-1.

References

[1] Luxand facesdk, 2013. 3

[2] Microsoft research face sdk beta, 2013. 3

[3] The talking face video, 2013. 4

[4] P. N. Belhumeur, D. W. Jacobs, D. J. Kriegman, and N. Kumar. Localizing parts of faces using a consensus of exemplars. In *CVPR*, 2011. 1, 3, 4

[5] X. Cao, Y. Wei, F. Wen, and J. Sun. Face alignment by explicit shape regression. In *CVPR*, 2012. 1, 4

[6] L. Chevallier, J.-R. Vigouroux, A. Goguey, and A. Ozerov. Facial landmarks localization estimation by cascaded boosted regression. In *International Conference on Computer Vision Theory and Applications (VISAPP)*, 2013. 2, 3, 4

[7] M. Dantone, J. Gall, G. Fanelli, and L. Van Gool. Real-time facial feature detection using conditional regression forests. In *CVPR*, 2012. 3

[8] P. Dollár, P. Welinder, and P. Perona. Cascaded pose regression. In *CVPR*, 2010. 3

[9] J. H. Friedman. Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, 2001. 2

[10] G. E. Hinton and R. R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 2006. 1

[11] O. Jesorsky, K. J. Kirchberg, and R. W. Frischholz. Robust face detection using the hausdorff distance. pages 90–95. Springer, 2001. 3

[12] Z. Kalal, K. Mikolajczyk, and J. Matas. Tracking-learning-detection. *TPAMI*, 2012. 2

[13] N. Markuš, M. Frljak, I. S. Pandžić, J. Ahlberg, and R. Forchheimer. A method for object detection based on pixel intensity comparisons. http://arxiv.org/abs/1305.4537v1, 2013. 2

[14] N. Markuš, M. Frljak, I. S. Pandžić, J. Ahlberg, and R. Forchheimer. Eye pupil localization with an ensemble of randomized trees. *Pattern Recognition*, 2014. 2

[15] E.-J. Ong, Y. Lan, B. Theobald, and R. Bowden. Robust facial feature tracking using selected multi-resolution linear predictors. *TPAMI*, 2011. 3

[16] M. Ozuysal, P. Fua, and V. Lepetit. Fast keypoint recognition in ten lines of code. In *CVPR*, 2007. 1, 2

[17] J. Shotton, A. Fitzgibbon, M. Cook, T. Sharp, M. Finocchio, R. Moore, A. Kipman, and A. Blake. Real-time human pose recognition in parts from single depth images. In *CVPR*, 2011. 2

[18] Y. Sun, X. Wang, and X. Tang. Deep convolutional network cascade for facial point detection. In *CVPR*, 2013. 1, 3, 4

[19] A. Torralba and A. A. Efros. Unbiased look at dataset bias. In *CVPR*, 2011. 4

[20] A. Vinciarellia, M. Pantic, and H. Bourlarda. Social signal processing: Survey of an emerging domain. *Image and Vision Computing*, 2009. 1

Figure 5: Some results obtained by our system on real-world images. Notice that we use more facial landmark point detectors than in the experimental section.